arXiv:1511.04803v2 [stat.ME] 17 Nov 2015

Performance and Interpretation of Classification Models

Yuan-chin Ivan Chang and Chih-Li Sung

Institute of Statistical Science, Academia Sinica, 128, Academia Road Section 2, Taipei Taiwan 11529

Abstract

Classification is a common statistical task in many areas. In order to ameliorate the performance of the existing methods, there are always some new classification procedures proposed. These procedures, especially those raised in the machine learning and data-mining literature, are usually complicated, and therefore extra effort is required to understand them and the impacts of individual variables in these procedures. However, in some applications, for example, pharmaceutical and medical related research, future developments and/or research plans will rely on the interpretation of the classification rule, such as the role of individual variables in a diagnostic rule/model. Hence, in these kinds of research, despite the optimal performance of the complicated models, the model with the balanced ease of interpretability and satisfactory performance is preferred. The complication of a classification rule might diminish its advantage in performance and become an obstacle to be used in those applications. In this paper, we study how to improve the classification performance, in terms of area under the receiver operating characteristic curve of a conventional logistic model, while retaining its ease of interpretation. The proposed method increases the sensitivity at the whole range of specificity and hence is especially useful when the performance in the high-specificity range of a receiver operating characteristic curve is of interest. Theoretical justification is presented, and numerical results using both simulated data and two real data sets are reported.

Keywords:
Area under curve; generalized additive model; logistic regression model; ROC curve

1. Introduction

Due to the predictive nature of classification models, the predictive power is usually the most important factor for assessing a classifier/diagnosis procedure. This is especially the case in the scenarios of disease screening or survey research, where an efficient classification rule that can effectively distinguish the diseased from normal subjects is preferred. However, we are also aware that the performance should not be the solo measure for assessing a classifier when the model interpretation and the impacts of factors are of interest. For

example, in some medical/pharmaceutical research, we want to learn the impacts of biomarkers in a classification rule. Because this information might also imply their associations with the studied disease, it could be useful for drug and therapy developments. Hence, in this kind of research, besides the diagnostic power, the ability of to identify useful biomarkers/variables or the ease of model interpretation is also an important criterion for evaluating a classification rule. In modern classification literature, we can easily find many powerful classification procedures, such as the support vector machine (see Cristianini and Shawe-Taylor, 2000), which can provide high classification accuracy. Yet, we also know that when these types of powerful and complicated classification rules are used, extra effort is required to interpret their rules/models, and to dig out the impact information of individual variables/biomarkers from them is even more difficult.

On the other hand, simple conventional statistical models, such as linear logistic models, can usually be easily interpreted; nevertheless, they might not have as good classification performance as those modern methods, especially when the association between diseases and variables is complicated. A nonlinear model can usually fit the observed data better than its linear model counterpart. However, can it also assure us good prediction performance in general classification problems? As far as we know, there is still a lack of theoretical justification in the literature to guarantee this. In fact, the inconsistency between the fitted model and the prediction-based model building has been reported in Tian et al. (2007). Of course, if the model is true, then fitting observations may better insure this model against bad predictions. However, in practical situations, the true model is usually unknown. Then, to overfit the observed data or training samples with a complicated model cannot guarantee good classification/prediction performance. From practical perspectives, to balance the classification performance and the ease of model interpretation is important, especially in the studies that focus on providing detailed model information for further research and developments. This motivates us to find a way to boost the performance of the conventional model and retain its ease of model interpretation.

The assessment measures play an important role in classification/diagnostic research (see Pepe, 2004), and using a different assessment measure commonly results in selecting a different classifier. Because the accuracy of a binary classifier can easily be affected by the selected cutting point, and because using training samples to decide a cutting point might easily cause overfitting, a cross-validation procedure is usually used to prevent this situation. However, the cross-validation method is computationally intensive. Thus, threshold independent measures, such as the receiver operating characteristic (ROC) curve, are usually preferred to accuracy (Pepe, 2004). The area under ROC curve (AUC) is a popularly used summarization statistic of the ROC curve, and in this study, we will use it as the performance measure.

For a binary classification problem, it has been proved in Eguchi and Copas (2002) that a logistic-type classification function will achieve the maximum AUC under certain conditions. Moreover, as stated in Hastie and Tibshirani (1986,

1990), the general additive model (GAM) is a method of approximating the true model and is a flexible statistical method that can be used to characterize and identify nonlinear regression effects. Many advantages of GAM have been intensively discussed. Among them, a commonly mentioned one is its ease of model interpretability (Hastie and Tibshirani, 1986, 1990; Wood, 2006). Hence, in this paper, a general additive logistic type of classification function is used to boost the classification performance of conventional models in terms of AUC. Theoretical justification is presented, and some suggestions for fitting a general additive logistic model are discussed.

The rest of this paper is organized as follows. We start with a brief review of the ROC curve, AUC, and GAM. Then, we will show the relations between the ROC curve and the logistic regression models, and how the AUC of a logistic model can be improved with a general additive logistic model. It is followed by the numerical results using both synthesized data and two real data sets. Fitting a GAM model relies on some iterative algorithms, and the different approaches used in the algorithms may affect the fitting results. Hence, the computational methods of GAM are compared, and some suggestions based on the empirical results are given in the Summary. The information on algorithms and the corresponding packages used here are described in the Appendix.

2. Review and Methodology

2.1. ROC cuver and AUC

The ROC curve and AUC are important statistical tools for evaluating binary classifiers (see Pepe, 2004). Let Z be the score of a binary classification function; then they can be briefly described as follows. Assume that the greater the value of Z, the higher the probability of a subject being diagnosed as a diseased one. Suppose $c \in R$ is a threshold and the subject with $Z > c$ will be classified as diseased. Then, the sensitivity and specificity for such a cutting point c are defined as $Se(c) = pr(Z > c)$ and $Sp(c) = pr(Z > c)$, respectively. It follows that the corresponding ROC curve is defined as a plot of $\{(1 - Sp(c), Se(c)) : -\infty < c < \infty\}$, and its AUC is then equal to $\int_0^1 Se(Sp^{-1}(t))dt$. That is, the plot of an ROC curve takes all possible cutting points into consideration, and the AUC of it can be viewed as a summarization of ROC curve. The detailed properties of the ROC curve and AUC can be easily found in the literature and textbooks, such as Pepe (2004).

2.2. General additive logistic models

Let $x_1, x_2, \ldots, x_p$ be predictors and $y \in \{0, 1\}$ be a binary response variable. A GAM has a form as

$$E(y|x_1, x_2, \ldots, x_p) = \alpha + f_1(x_1) + f_2(x_2) + \ldots + f_p(x_p), \tag{1}$$

where f_j's are some smooth functions. Thus, it can be viewed as an additive extension of the family of generalized linear models (Hastie and Tibshirani,

1990). Each smooth function in (1) can be fitted using a scatterplot smoother, such as a cubic smoothing spline or kernel smoother. As a special case of the generalized linear models, a general additive logistic model, under the GAM framework, is formulated as

$$\log\left(\frac{p(x)}{1-p(x)}\right) = \alpha + f_1(x_1) + f_2(x_2) + \ldots + f_p(x_p), \tag{2}$$

where $p(x) = P(y = 1|x)$ and f_j's are some unknown smooth functions. The smooth functions of GAMs are estimated based on the observations with some iterative procedures. Due to the additive assumption, these smooth functions can always provide the relations between the variables and the disease. The non-parametric form of f_j gives us some flexibility for model fitting, and the additive form of GAM retains much of the ability of model interpretation (Hastie et al., 2001). Besides these properties, there are more classification-related advantages of the additive logistic models discussed in Friedman et al. (2000).

In many medical diagnostic situations, to learn the relationship between the sensitivity and specificity of classifiers is important. An ROC curve and its AUC can provide more information before a threshold is determined. Thus, despite the accuracy of a classification rule, they are commonly used measures in these kinds of research. Below, we show how the classification power of a conventional logistic model can easily be improved by an additive logistic model counterpart with only a little extra computational effort.

2.3. AUC of additive logistic models

Let y be a binary response variable as before. Suppose that $y = 0$ or 1 are labels of two classes in a binary classification problem. Let $x \in R^p$ be the vector of measurements of a subject and $g_0(x)$ and $g_1(x)$ be the conditional probability density functions of $x \in R^p$, given $y = 0$ and $y = 1$, respectively. Let π_0 and π_1, with $\pi_0 + \pi_1 = 1$, denote the probabilities of $y = 0$ and $y = 1$, accordingly. That is, π_0 and π_1 are the proportions of two classes in the whole population. It follows that for a given x, the log-odd ratio of two classes, say $\lambda(x)$, can be written as a function of x:

$$\lambda(x) = \log\frac{p(x)}{1-p(x)} = \log\frac{\pi_1}{\pi_0} + \log\frac{g_1(x)}{g_0(x)}. \tag{3}$$

As stated before, an ROC curve is a plot true positive rate (sensitivity) against the false positive rate (1-specificity). If we treat the false positive and false negative as Type I and Type II errors, respectively, then an ROC curve can also be viewed as a plot of Type II errors versus Type I errors at different critical values. Moreover, from the Neyman-Pearson lemma, we have that for a given Type I error, the likelihood ratio test is the most powerful test for testing the null hypothesis $y = 0$ versus the alternative hypothesis $y = 1$. Hence, under this hypotheses testing framework, and using the Neyman-Pearson lemma, it is proved that $\lambda(x)$ defined in (3) is a statistic that maximizes the height of the ROC curve at all levels of specificity. That is, if $\lambda(x)$ is used as a classification

function, then its corresponding ROC curve will dominate the ROC curves of all other binary classification functions and therefore has the maximum AUC. Based on this fact, Eguchi and Copas (2002) showed that the maximum AUC can only be achieved when the true model for the density ratio, $g_1(x)/g_0(x)$, is known. However, both $g_0(x)$ and $g_1(x)$ are usually unknown in practice, and so is $\lambda(x)$. Thus, among all possible classification functions based on the observations, the one that can best approximate $\lambda(x)$ will have the largest maximum AUC. It follows that when a conventional logistic model is a true model, then the classification function based on it will have the maximum AUC. Therefore, in this case there is no room for further improvement. Otherwise, if the prediction power is the only concern, then one can always find a more complicated model for approximating the likelihood ratio in order to have better classification performance. On the other hand, if the final model is too difficult to be interpreted, then it provides little information about the relation between responses and the variables of interest. Thus, these types of models will not be useful for identifying the impacts of individual variables on the response. However, in many situations, to identify the key variables is also an important task. For example, in pharmaceutical and medical research, this information is highly valued for developing new drugs or therapies. It is well-known that a general additive model will usually fit data well. In addition, authors suggest that to use an additive logistic model to fit data is a good strategy in practice, as it will retain some ease of model interpretation (Friedman et al., 2000). However, it is also pointed out in the literature that a better-fitting model cannot guarantee a better prediction (Tian et al., 2007) in a general sense, and as far as we know, there is still no theoretical proof that a better-fitting model will also have a larger AUC. In below, we show that a general additive model will have a larger AUC than its linear logistic model counterpart will.

Proof. Let $\lambda(x)$ be the true log-odds ratio as before. If a linear logistic model is used to approximate the log odds-ratio, then we assume that $\log[p(x)/1-p(x)] = \beta' x$, where β is the unknown regression parameters to be estimated. Similarly, if an additive logistic model is used, then we set $\log[p(x)/1-p(x)] = \Sigma_i f_i$, where f_i's are some smooth functions as in (2). Let u, u' and u'' be three constants and define $A = \{x : \lambda(x) > u''\}$, $B = \{x : \Sigma_i f_i > u'\}$, $C = \{x : \beta' x > u\}$ as the corresponding critical regions of the likelihood ratio tests based on $\lambda(x)$, $\Sigma_i f_i$ and $\beta' x$, respectively, such that

$$\int_A g_0(x)dx = \int_B g_0(x)dx = \int_C g_0(x)dx. \tag{4}$$

That is, all three test statistics, $\lambda(x)$, $\Sigma_i f_i$ and $\beta' x$, have the same Type I error. Then, by the Neyman-Pearson lemma, we have

$$\int_A g_1(x)dx - \int_C g_1(x)dx \geq 0, \tag{5}$$

and the equality in (5) holds only when the linear logistic model is the true model for the log odds-ratio. If this linear logistic model is not the true model,

then there is an $\epsilon_1 > 0$ such that

$$\int_A g_1(x)dx - \int_C g_1(x)dx \geq \epsilon_1, \tag{6}$$

Suppose that x_j's are bounded; then the estimates of f_i's, say f_i^*'s, are chosen to minimize $E[(\lambda(x)-\Sigma_i f_i^*)^2]$ subject to the constraint $E(f_i^*) = 0$. Then, under a common smoothness assumption on f_i^*, Stone (Stone, 1985) showed that the rate of convergence for spline estimates of $\Sigma_i f_i^*$ is $n^{-(2p)/(2p+1)}$, where n is the sample size and p is a measure of the smoothness of f. Thus, $E[(\Sigma_i f_i(X) - \lambda(X))^2]$ will converge to 0 as the sample size becomes large. Because (4) holds, it implies that for any given $\epsilon_2 \geq 0$, there is a large enough n such that

$$\epsilon_2 \geq \int_A g_1(x)dx - \int_B g_1(x)dx \geq 0. \tag{7}$$

Because ϵ_2 can be arbitrarily small as n goes to infinity, equations (5) and (7) together imply that as the sample size n is large enough,

$$\int_B g_1(x)dx - \int_C g_1(x)dx \geq \epsilon_1 - \epsilon_2 > 0. \tag{8}$$

It implies that for a large enough n, we have

$$\int_B g_1(x)dx - \int_C g_1(x)dx > 0.$$

That is, when the sample size is large enough, there is a spline estimate, $\Sigma_i f_i$, of $\lambda(x)$ such that

$$\int_A g_1(x)dx \geq \int_B g_1(x)dx > \int_C g_1(x)dx. \tag{9}$$

Let $\delta(u)$ be a function of u, $-\infty < \delta(u) < \infty$ such that $\text{pr}\{\Sigma_i f_i > u|y = 0\} = \text{pr}\{\beta' x > u + \delta(u)|y = 0\}$. Therefore, $\text{pr}\{\Sigma_i f_i > u|y = 1\} \geq \text{pr}\{\beta' x > u + \delta(u)|y = 1\}$ as $\text{pr}\{\Sigma_i f_i > u|y = 0\} = \text{pr}\{\beta' x > u + \delta(u)|y = 0\}$. Let AUC_{GAM} denote the AUC of the general additive logistic model based on $\Sigma_i f_i$, and let AUC_{GLM} denote the AUC of the linear logistic regression model based on $\beta' x$. Then, the rest of the proof follows via similar arguments of that in Eguchi and Copas (2002). That is, it is shown that

$$\begin{aligned}
\text{AUC}_{\text{GAM}} &= \int_{u=+\infty}^{-\infty} \text{pr}\{\Sigma_i f_i > u|y = 1\} d(\text{pr}\{\Sigma_i f_i > u|y = 0\}) \\
&= \int_{u=+\infty}^{-\infty} \text{pr}\{\Sigma_i f_i > u|y = 1\} d(\text{pr}\{\beta' x > u + \delta(u)|y = 0\}) \\
&\geq \int_{u=+\infty}^{-\infty} \text{pr}\{\beta' x > u + \delta(u)|y = 1\} d(\text{pr}\{\beta' x > u + \delta(u)|y = 0\}) = \text{AUC}_{\text{GLM}}.
\end{aligned}$$

This completes that proof that as n is large, the general additive logistic model based on $\Sigma_i f_i$ has a larger AUC than the linear logistic model with $\beta' x$.

The above result says that unless the linear logistic model above is the true model, which leaves no room for improvement, there is always a general additive logistic model that has a larger AUC than that of the original linear logistic model when the sample size is large. Moreover, because (9) holds for all classification thresholds, it implies that the general additive logistic model has larger sensitivity than its corresponding linear logistic model does at any specificity level. That is, the general additive logistic model actually also has a larger partial AUC than the linear logistic model does at any given range of specificity. This is a very useful feature, especially in rare disease diagnosis or other similar situations where the high specificity of a classification model is a must.

3. Empirical Study

We report the numerical results using both synthesized and real data sets. There are several numerical algorithms available in the R platform that can be used to fit a general additive logistic model with different options. According to the computational approaches adopted in these algorithms, there are three categories:

(a) back-fitting method (see also Hastie and Tibshirani (1990), Chapter 4),

(b) simultaneous estimation method (estimating all components with optimization in smoothing parameter space (Marx and Eilers, 1998; Wood, 2004), and

(c) likelihood-based boosting method (Tutz and Binder, 2006).

The computational efficiency depends on many factors such as the dimensionality of explanatory variables of the data set, sample size and so on. Basically, we found that none of these three types of methods can dominate the other two in all situations. Hence, in order to provide useful information for practitioners to choose a suitable one for their needs, the advantages and disadvantages of using these methods in different situations are discussed. The results of fitting a linear logistic regression with a conventional method are used as baseline models, and the results of the general additive logistic models using different fitting algorithms will be compared with these results. (The variable selection is not the focus in this paper; the variable selection schemes will be used only when there is such an option available in the corresponding packages.) The following abbreviations denote the algorithms and the options used in our numerical studies.

1. *GLM*: Fitting a *linear logistic regression model* with all predictors
2. *GLM_step*: Fitting a *linear logistic regression model* with a backward elimination method for variable selection

3. *Backfitting*: Fitting a *general additive logistic model* using a back-fitting algorithm
4. *Backfitting_step*: Fitting a *general additive logistic model* using a back-fitting algorithm with the AIC option for stepwise selection of model components and degrees of freedom
5. *mgcv*: Simultaneously estimating all components with smoothing parameter optimization when fitting a *general additive logistic model*
6. *mgcv_step*: Similar to *mgcv* with automatic selection of smoothing parameter using the AIC criterion
7. *GAMBoost*: A boosting procedure with the number of boosting steps selected with the AIC criterion

Because fitting a GAM relies on the iterated algorithms and some parameter-setting in the algorithms need to be set in advance, there are many suggestions for the parameter setting in the algorithms in Wood (2008); Hastie et al. (2001); Tutz and Binder (2006). Here, we follow the suggestions of Tutz and Binder (2006) by apply a scaling factor 1.4 to the effective degrees of freedom of *mgcv*, and we multiply the degrees of freedom of both *GLM* and *backfitting* methods by 1.4. For the *GAMBoost* has a variable selection option during its boosting procedure, so we use the default variable selection scheme in *GAMBoost*. All numerical studies are conducted using the R language on a 64-bit Linux PC cluster with Xeon Octa-Core and Hyper-Threading E5-2670 2.6 GHz CPU with 256GB of RAM. The technical details of the packages used here, please refer to their original manuals in R.

In addition to AUCs, the computational times of using these methods in different model-fitting scenarios in our numerical studies are reported. The fitting methods are described in the Appendix. We just use the variable selection option when it is available in a package; however, we will not discuss much about it. For the details of the variable selection option, please refer to their original papers, which can be found in the documents of the corresponding packages. Instead, we emphasize the improvement of AUCs when a general additive logistic type classification function is used and its interpretation.

3.1. Synthesized data

We consider two different dimensions ($d = 5$ and $d = 10$) of models. To simulate different kinds of model complexities, we use two sets of effective variables as described in equations (10) and (11) below:

$$\text{Set 1: } g_1 = x, \quad g_2 = 2x^2, \quad g_3 = \sin(5x), \tag{10}$$

$$\text{Set 2: } g_1 = 2(-x^3 + 1), \quad g_2 = 3\exp(-5x^2), \quad g_3 = 4\log(1 + x^2). \tag{11}$$

The others variables, $x_4, \ldots, x_d$ of the model are generated from a uniform distribution $U[-1, 1]$. That is, in all cases, only Variable Set 1 and Set 2 above have contributions to their corresponding responses, and the other variables $x_4, \ldots, x_d$ will have no contributions at all. In our studies, the true log-odds function of models with Set 1 and Set 2 variables have the same form below:

$$\eta = 3(-0.7 + g_1(x_1) + g_2(x_3) + g_3(x_5)).$$

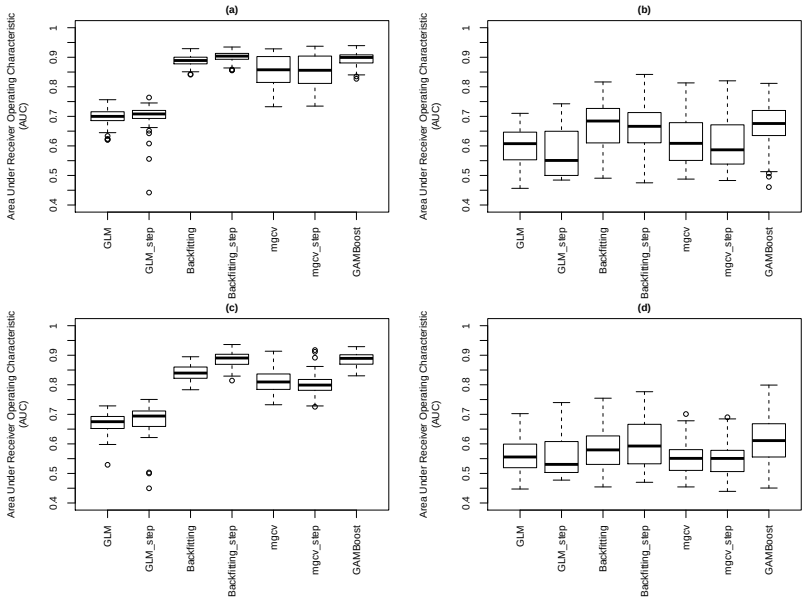

Figure 1: Box-plots of AUCs with 100 repetitions in four difference scenarios: (a) $d = 5$ with Equation (10), (b) $d = 5$ with Equation (11), (c) $d = 10$ with Equation (10) and (d) $d = 10$ with Equation (11).

The true probability vector $p = (p_1, ..., p_n)'$ is calculated with $p_i = \exp(\eta_i)/(1 + \exp(\eta_i))$ for each i, and the binary response vector $y = (y_1, ..., y_n)'$ is generated using a Bernoulli distribution with success probability p_i for each $i = 1, \ldots, n$. There are 100 repetitions for each case. For each run, we separately generated 100 training samples and 1000 testing samples.

Figure 1 shows the box-plots of AUCs based on 100 replications for different models using different fitting methods described before. For the model with Set 1 as its effective variables, all fitting methods for the general additive model have larger AUCs than the conventional linear logistic models do (with and without variable selection option). For the model with the effective variable Set 2, the advantage of the general additive logistic model is not as significant as the previous case, which is due to the complexity of the true model such that both the linear and general additive models cannot approximate the true model well. However, the additive logistic model using *backfitting* and *GAMBoost* still outperforms the linear logistic model.

Figure 2 are plots of the average ROC curves in the left-upper corner of the unit-square, $[0, 1] \times [0, 1]$. We can clearly see that there are two groups of ROC curves in each picture, and the ROC curves of the general additive model, with different fitting methods, are always higher than that of logistic models (GLM and GLM_step). The general additive model fitted with different

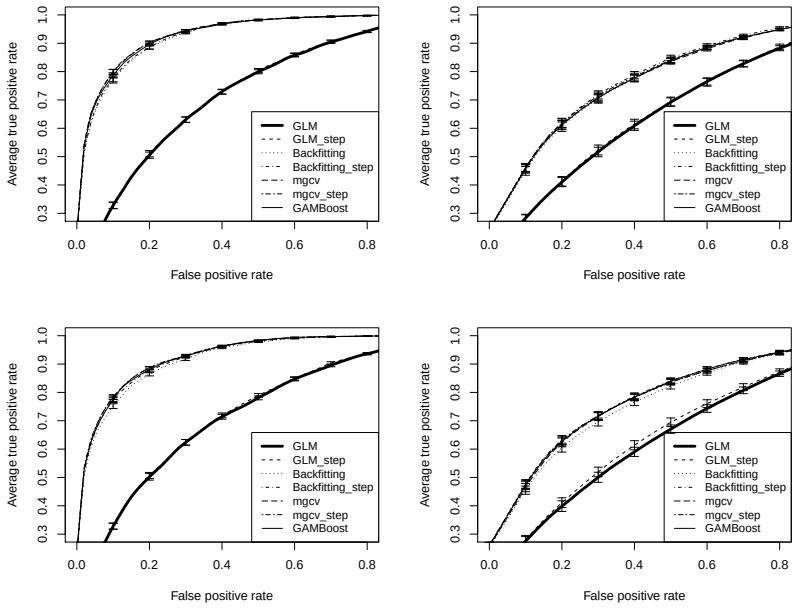

Figure 2: The average ROC curve at the left-upper corner of the $[0,1] \times [0,1]$ unit square for each method. The curve of each method is based on 100 repetitions, and data are generated from 4 different situations: (a) $d = 5$ with Equation (10), (b) $d = 5$ with Equation (11), (c) $d = 10$ with Equation (10) and (d) $d = 10$ with Equation (11). The vertical bars in these plots indicate the 95% point-wise confidence intervals for each curve.

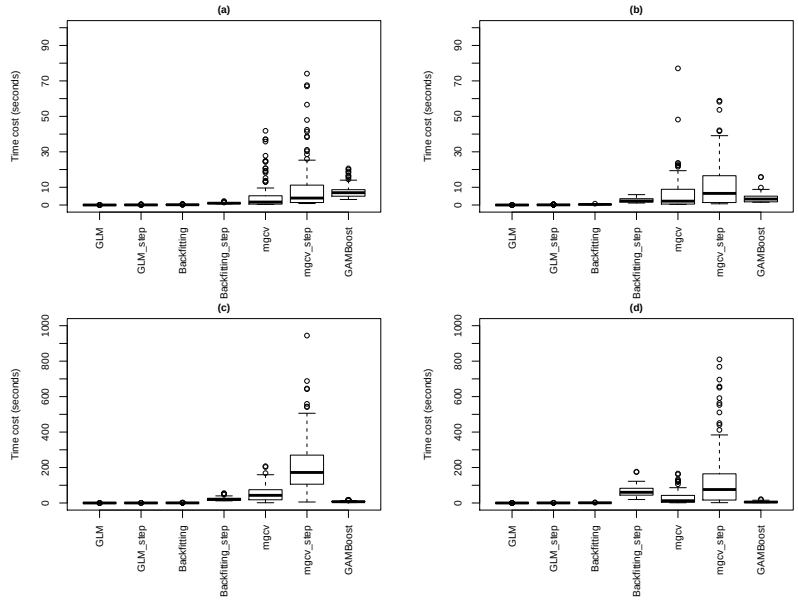

Figure 3: The CPU times for fitting models using different methods, based on 100 repetitions, with data generated from (a) $d = 5$ with Equation (10), (b) $d = 5$ with Equation (11), (c) $d = 10$ with Equation (10) and (d) $d = 10$ with Equation (11).

algorithms is essentially the same model unless their default variable selection option is used, which may introduce some variation due to their variable selection schemes. (Note that GLM denotes the result of fitting a conventional linear logistic model, and GLM_step denotes the result of a similar fitting procedure but with a stepwise variable selection option.)

Figure 3 shows the box-plots of computational times of different methods. When $d = 5$, the *backfitting* method uses the shortest time among all methods in two models. When $d = 10$, the *GAMBoost* method becomes very competitive, especially when the model is complicated as in (11) (see Figure 3(d)). Among different fitting methods for the general additive logistic model, the *backfitting* and *GAMBoost* are more computationally efficient than are the *mgcv* methods with and without step-wise variable selection. Based on the simulation results, the *GAMBoost backfitting* methods are "cost-effective" in terms of their gains in AUCs and the little cost of extra computational time used compared to that of fitting a conventional linear logistic model with *GLM*-based algorithms. There are more discussions and comparisons of the computational-related issues for algorithms for fitting a logistic model and GAM that can be found in Minka (2007); Binder and Tutz (2008).

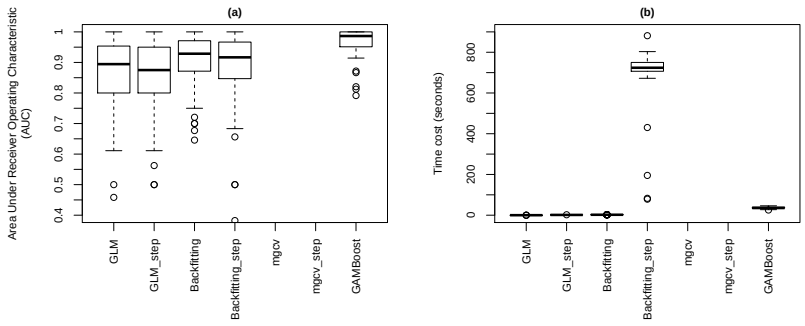

Figure 4: (a) AUC and (b) computational cost time measured by each fitting method with 100 resampling data from parkinson's data set.

3.2. Real data examples

We also apply the proposed method and algorithms discussed above to two real data sets, Parkinson's disease (Little et al., 2007) and cardiotocography (Bache and Lichman, 2013), for illustration purposes. We compared the performances of methods using a general additive logistic model to the results obtained via fitting a linear logistic model. The main purpose of using these two data sets is to assess the improvement in AUCs when replacing a linear logistic with a general additive logistic model. Hence, for the details of these two data sets, please refer to their original papers (Little et al., 2007; Bache and Lichman, 2013). Brief information about these two data sets is summarized below.

Parkinson's disease data set. The original purpose of this data set is to find a classification rule to discriminate healthy people from the Parkinson's disease (PD) patients based on some biomedical voice measurements. There are 31 persons in this study, and among them, 23 are patients. The study includes 195 voice recordings from these 31 individuals with different numbers of replications of each individual. There are 23 variables for each recoding. The binary code indicates the voice recording is from a diseased or a normal subject. The other 22 predictive variables consist of the records of several voice measurements, such as average vocal fundamental frequency, variation in amplitude and so on. For illustration purposes, we treat 195 voice recordings as 195 individual samples. Thus, among these 195 voices, there are 147 voice recordings from PD subjects and 48 from healthy subjects. That is, we treat it as a data set with 195 samples (147 positive and 48 negative cases) with 22 variables, and the target becomes to build a classification rule to distinguish these two types of recordings.

We follow the common training-test framework in machine learning and classification literature. Each time, 90% of samples are randomly selected from these two groups for the model fitting (training), and the remaining 10% of each of them are used as testing samples. This procedure is then repeated 100 times. Because of the number of predictive variables is relatively large comparing to the subject size we have, *mgcv* and *mgcv_step* methods cannot converge in a

reasonable computing time. Hence, we do not include the results of these two algorithms in this example. Figures 4 (a) and (b) are the box-plots of AUCs and the computational times of all other methods based on 100 runs, respectively. In this example, all methods, including *GLM* and *GLM_step*, perform well in terms of their AUC averages. The *G*AMBoost has the largest average AUC with the smallest variation among these methods. The linear logistic models (with and without a stepwise variable selection option) have smaller means of AUCs and larger variances than others do. Most of the algorithms are computationally efficient in this case, except *backfitting_step*. This is due to the stepwise variable selection procedure in *backfitting_step*. Without stepwise option, the *backfitting* is as efficient as others. The differences of the computational times of *GLM*, *GLM_step* and *backfitting* are not statistically significant. In fact, the CPU time used in *backfitting* is only slightly longer than that of *GLM* and *GLM_step*. Thus, from both the improvement of AUCs and the corresponding computational time view points, *backfitting* and *GAMBoost* are very "cost-effective" in this case.

The advantage of the *backfitting* and *GAMBoost* can also be seen from the plot of the average ROC curves in Figure 5. In order to have a clearer picture, we only show the left-upper $[0, 0.5] \times [0.5, 1]$ corner of the unit square. Each curve plotted here is based on 100 replications, and the vertical bars on each curve are the corresponding 95% point-wise confidence intervals at the given false positive rates. It is clearly shown that the ROC curve of *GAMBoost* is significantly higher than other ROC curves in this range of false positive rate, and all ROC curves from fitting a general additive logistic model, with different fitting algorithms, are higher than those of GLM methods. This result confirms that the recommended method is useful even when only a specific range of false positive rate is of interest. (Note that the ROC curves of *GLM* and *GLM_step* are very close and thus cannot be distinguished in this figure.)

Figure 6 shows the plots of estimated smooth functions and their corresponding partial residual plots for the selected predictors obtained from the *backfitting_step* method. Because of the additive model assumption, we can look at one variable at a time by keeping the others fixed. Thus, the interpretation of the relations between predictors and response variable is similar to that in a linear logistic model (see Hastie and Tibshirani, 1990). From this figure, we can see the impacts of the individual variables on the odds-ratio. For example, the left-upper corner is a plot of the curve of the estimated association function for MDVP.Fo.Hz, and it shows the functional relation between the log odds-ratio (response variable) and this variable. When the values of all other variables are fixed, this plot shows that MDVP.Fo.Hz and the response is positively associated at the beginning and then turns to negatively associated as the value of the variable becomes large. Compared to that in the association plot for MDVP.Flo.Hz at the first row and second column, it clearly shows that these two variables have different associations with the responses. Because the curve of MDVP.Flo.Hz stays near 0, it indicates that the association of MDVP.Flo.Hz with the responses is not as strong as MDVP.Fo.Hz. In addition, the picture at the second row and forth column shows that the estimated function of variable

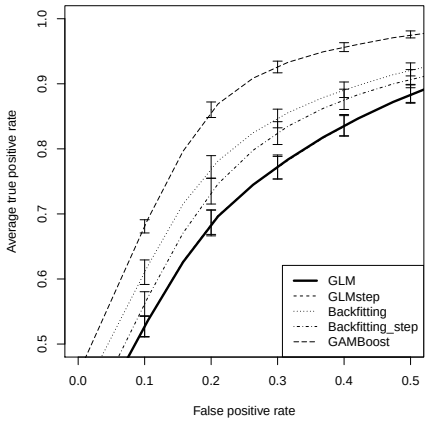

Figure 5: Average ROC curves with confidence bars are obtained by each fitting method from Parkinson's data set with 100 replications.

MDVP.APQ is almost flat, and the ± 2 standard deviation curves of it cover 0 at most values of *MDVP.APQ* in this picture, except for the values around 0.02. This suggests that this variable might have just a little impact on the (log) odds-ratio (see Marra and Wood, 2012; Binder and Tutz, 2008). (Please note that we just use this as an example for illustration purposes and no disease-related inference is made here.) Thus, we can clearly find the impacts of individual variables on the responses through these plots as well as the certain levels of the interpretation ability of variables. It is clear that we cannot obtain this kind of nonlinear relation using a linear logistic model. The way of interpretation is similar to that of a linear logistic model except that they are in functional form instead of slopes. More discussions about the interpretation of GAMs can be easily found in the literature such as Tutz and Binder (2006); Wood (2006) and textbooks about GAM as Marx and Eilers (1998); Hastie and Tibshirani (1990); Hastie et al. (2001).

Cardiotocography data set. The cardiotocography data set used here consists of 2126 fetal cardiotocograms (CTGs), where 19 diagnostic features were measured, including fetal heart rate and uterine contraction features. Based on the judgements of some expert obstetricians, these 2126 subjects are classified as three fetal states: normal, suspect, and pathologic. Because we consider only two-class classification problems in this work, we combine two neighbored classes into one class. Thus, they becomes two binary classification problems: (1) {**normal and suspect**} versus {**pathologic**} and (2) {**normal**} versus {**suspect and pathologic**}. For demonstration, they are treated separately as two different binary classification problems.

Besides the contents, a major difference between these two real data examples is that the cardiotocography data set has more observations and less

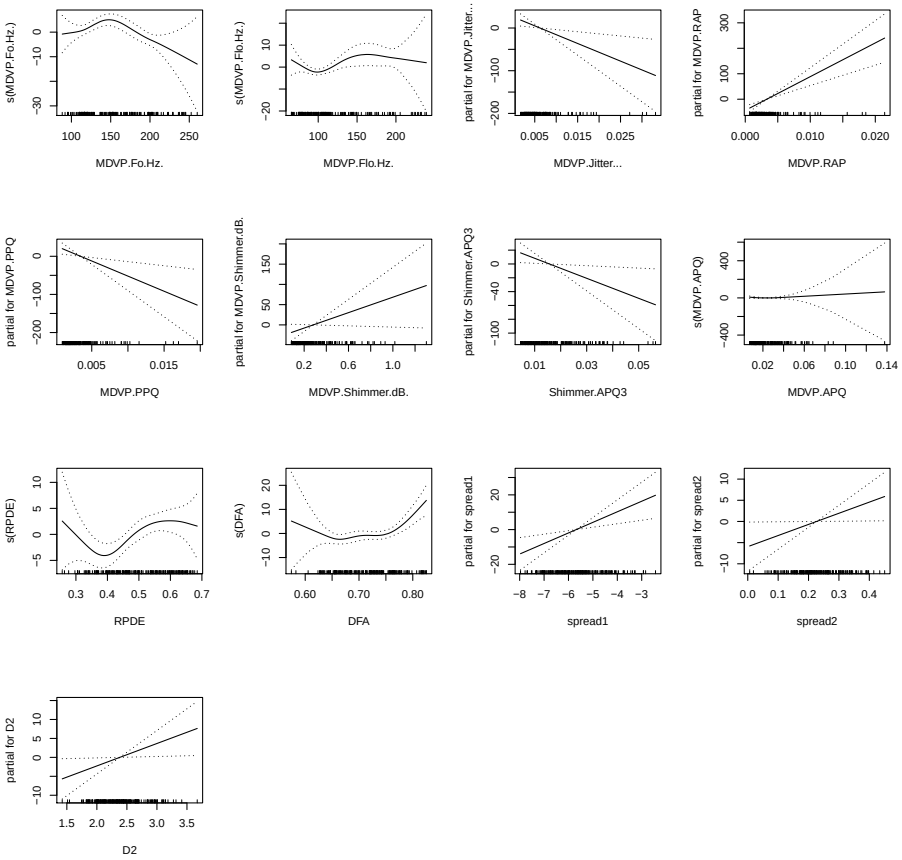

Figure 6: The plots of the estimated functions of 13 predictors selected in the *backfitting_step*, where x-axis is the value of each predictor, the solid line is the estimated functions and the dot-lines represents $\pm$ 2 standard-error curves obtained from the fitting algorithm.

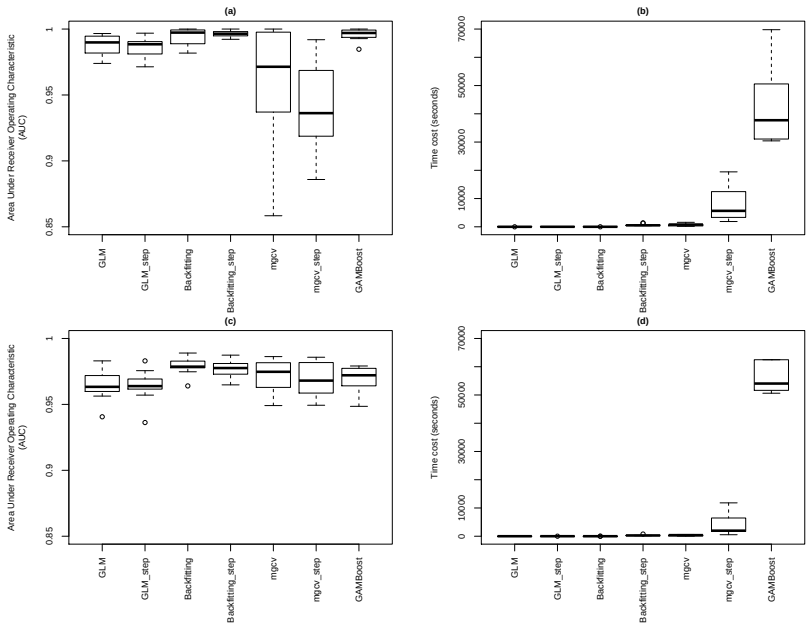

Figure 7: (a) AUC and (b) computational cost time measured by each fitting method with 10 resampling data from cardiotocography data set if label "suspect" belongs to "normal." (c) AUC and (d) computational cost time measured by each fitting method with 10 resampling data from cardiotocography data set if label "suspect" belongs to "pathologic."

predictors, which will affect the computational efficiency of some algorithms. The results reported here are based on a smilier re-sampling plan as before. Both *mgcv* and *mgcv_step* algorithms will now converge in a reasonable time, so their results are included. In this example, we omit the plots of the estimated functions and only focus on AUCs and related model fitting issues that are useful for practitioners to choose suitable methods for their own studies.

Figures 7 (a) and (c) are the box-plots of AUCs for the cases of {normal and suspect} versus {pathologic} and {normal} versus {suspect and pathologic}, respectively. Figures 7 (b) and (d) show their corresponding computational times of all methods. Although *mgcv* and *mgcv_step* converge in a reasonable time in this example, they are still not as stable as other GAM-based counterparts are in terms of variation of AUCs. The *GAMBoost* algorithm takes a much longer time than others in this case. (Note that in our study, *GAMBoost* takes more than eight hours for one run, while the *backfitting* and *backfitting_step* take only less than five minutes.) This phenomenon is due to the large sample size of the cardiotocography data set. This result is consistent with our simulations and the comments reported in the literature. Besides *mgcv* and *mgcv_step*, all other GAM-based fitting algorithms perform better than *GLM* and *GLM_step* do in terms of their means of AUCs and variances. Hence, when the sample size is large and the number of predictive variables is relatively small, the *backfitting* and *backfitting_step* algorithms are recommended.

Figure 8 shows the ROC curves at the left-upper corner of the unit square, and Figure 8 (a) is {normal and suspect} versus {pathologic} case, and 8 (b) is for {normal} versus {suspect and pathologic} case. From Figure 8 (a), we can see that the ROC curves of *Backfittting*, *Backfitting_step* and *GAMBoost* are three methods on the top; meanwhile, then the two curves of *GLM*-based methods, and the curves of *mvcv*-based methods are the worst two in this case. Figure 8 (b) shows that ROC curves of the {normal} versus {suspect and pathologic} case. Here, the performances of all methods are very close. However, we still find that all *GAM*-based fitting algorithms are significantly better than the two *GLM*-based methods in the low false positive rate range (i.e., less than 0.15 in this plot). In fact, in both binary classification cases of this data set, using a general additive logistic model increases the sensitivity in the range of the low false positive rate. This range is usually important when a high-specificity diagnostic rule is preferred. Hence, this is also an advantage of using the recommended method.

4. Summary

In this paper, the conventional logistic classification function can be further improved in terms of AUC. We show that both numerically and asymptotically, the spline-based generalized additive logistic model will have a larger AUC than a conventional linear logistic model does. We also demonstrate that the classification performance of a conventional linear logistic regression model can be easily improved with a general additive logistic model without sacrificing much of its ease of interpretation using numerical results using some synthesized data

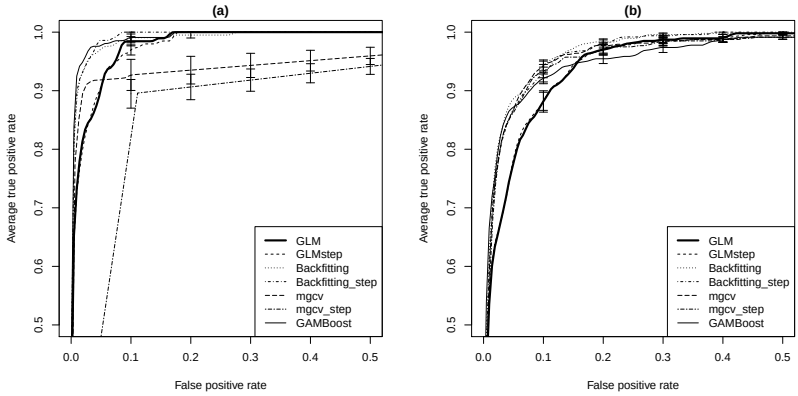

Figure 8: Average ROC curves with 95% confidence bars of each fitting method based on 100 resampling replications using cardiotocography data set for (a) {**normal and suspect**} versus {**pathologic**} and (b) {**normal**} versus {**suspect and pathologic**}.

sets and two real data sets. Several numerical algorithms are used to fit a GAM, and their pros and cons are discussed based on the numerical results in our studies. We found that the *GAMBoost* method outperforms other methods when with a moderate sample size and variable length. When the sample size becomes large, the *backfitting* algorithms are more computationally efficient compared to the CPU times used for fitting a conventional linear logistic model. It is worth mentioning that the most of the algorithms used here are "parallel-computing friendly." (Please see their manual for further details). That is, these algorithms can be used in some parallel computing environment, such as a PC-cluster or just a multi-core CPU, which is now very popular in desktop or portable computers, with little programing effort. Hence, the proposed method is useful for improving the diagnostic performance without diminishing the ability of the interpretation of the final prediction model, and it is "cost-effective" in terms of the computational time and the gain in AUC. Moreover, the ROC curve of the additive logistic model is higher than that of its linear logistic model counterpart in the whole range of specificity. It implies that the partial area under ROC curve in the high-specificity range is always improved, which is an important and useful feature in many medical diagnosis situations, and it is especially beneficial for applications where the high specificity is a must.

Appendix A: Algorithms for fitting general additive models

Backfitting algorithm. The backfitting algorithm employs a concept of conditional expectation; that is, if GAM is correct, then for any k, $E[Y - \alpha - \sum_{j \neq k} f_j(X_j)|X_k] = f_k(X_k)$. Following (1), the backfitting algorithm is stated below;

(i) Initial step: $\hat{\alpha} = N^{-1} \sum_1^N y_i, \hat{f}_j \equiv 0$ for $j = 1, \ldots, p$.

(ii) Cycle: $j = 1, 2, \ldots, p, \ldots, 1, 2, \ldots, p, \ldots,$

$$\hat{f}_j \leftarrow S_j[\{y_i - \hat{\alpha} - \sum_{k \neq j} \hat{f}_k(x_{ik})\}_1^N],$$

$$\hat{f} \leftarrow \hat{f} - \frac{1}{N} \sum_{i=1}^{N} \hat{f}_j(x_{ij}),$$

where $S_j(\cdot)$ denotes a smoothing spline for a target against the predictor x_j.

(iii) Continue Step (ii) until individual functions do not change.

For a generalized additive logistic model as in (2), Hastie and Tibshirani proposed a local-scoring algorithm (see Chapter 4 of their book), which is a maximum likelihood approach under a linear logistic regression model setup and a Newton-Raphson iterative type of method. The local-scoring algorithm has the following steps:

(i) Initial step: $\hat{\alpha} = \log(\bar{y}/(1-\bar{y}))$, $\hat{f}_j \equiv 0$ for $j = 1, \ldots, p$.

(ii) Cycle:

(a) Construct a new response:

$$z_i = \hat{\alpha} + \sum_{j=1}^{p} \hat{f}_j(x_{ij}) + \frac{y_i - \hat{p}_i}{\hat{p}_i(1-\hat{p}_i)},$$

where $\hat{p}_i = \exp\{\hat{\alpha} + \sum \hat{f}_j(x_{ij})\}/[1 + \hat{\alpha} + \exp\{\sum \hat{f}_j(x_{ij})\}]$.

(b) Construct weights: $w_i = \hat{p}(1-\hat{p})$.

(c) Estimate new $\hat{\alpha}$ and $\hat{f}_j$ for any j by fitting a weighted additive model to z_i with weights w_i via a backfitting algorithm.

(iii) Continue Step (ii) until the individual functions do not change.

(Note that the GAM package in R can perform both the backfitting and local scoring algorithms.)

Simultaneous estimation of all components with optimization in a smoothing parameter space. Both (1) and (2) can be modeled using the B-spline method: $f_j(x) = \sum_{i=1}^{K} \beta_{ij} B_{i,m}(x)$, where $B_{i,m}$ is the ith B-spline basis function of order m for the $K + m + 2$ evenly spaced knots, and β_{ij} is an unknown coefficient of $B_{i,m}$ for f_j. To prevent overfitting and have a better prediction, Marx and Eilers replaced the basic B-splines with a penalized B-splines, called P-splines (see Marx and Eilers, 1998). Thus, when estimating the unknown β_{ij} by maximizing a likelihood function, the P-splines will attach a different penalty on the adjacent coefficients of the B-splines to guarantee their smoothness. Therefore, to fit a GAM becomes to maximize the penalized likelihood below:

$$l^\star = l(\mathbf{y}; \boldsymbol{\beta}) - \frac{1}{2} \sum_{j=1}^{p} \lambda_j \boldsymbol{\beta}_j^T \mathbf{P} \boldsymbol{\beta}_j, \tag{A.12}$$

where λ, for all j, are smoothing parameters, $\mathbf{P}$ is the penalty matrix, and $l(\mathbf{y};\boldsymbol{\beta})$ is the log-likelihood of GAM. For a generalized additive logistic model as in (2), the log-likelihood with a penalized first order differences becomes

$$l^{\star} = l(\mathbf{y};\boldsymbol{\beta}) - \frac{1}{2}\sum_{j=1}^{p}\left\{\lambda_j \sum_{i=1}^{K-1}(\beta_{i+1,j} - \beta_{i,j})^2\right\},$$

where $l(\mathbf{y};\boldsymbol{\beta}) = \sum_{i=1}^{n}\{y_i \log p_i + (1-y_i)\log(1-p_i)\}$. Wood (2004) provided a stable and efficient approach to estimate all functions, simultaneously. This feature is available in the *mgcv* package in R.

Likelihood-based boosting method. Its basic idea of the boosting algorithm is to combine some "weaker classifiers" to produce a more powerful classifier that can reduce misclassification error using an iterative adaptive weighting scheme on observations. Tutz and Binder (2006) use it as a way of fitting an additive expansion in a set of elementary basis functions by minimizing a loss function, such as a likelihood-based loss function. They proposed the GAMBoost procedure to fit GAM with this boosting technology via maximizing the log-likelihood (A.12). In the statements below, we also use similar notations as that in the backfitting. The GAMBoost algorithm for fitting a generalized additive logistic model (2) is stated as follows:

(i) Definition: $\eta_i = \alpha + f_1(x_{i1}) + \ldots + f(x_{in})$, $\mu_i = h(\eta_i) = \exp(\eta_i)/[1+\exp(\eta_i)]$.

(ii) Initialize: $\hat{\alpha} = \log(\bar{y}/(1-\bar{y}))$, $\hat{f}_j = \hat{f}_j^{(0)} \equiv 0$ for $j = 1,\ldots,p$, yielding $\hat{\boldsymbol{\eta}}^{(0)T} = (\hat{\eta}_1^{(0)},\ldots,\hat{\eta}_n^{(0)}) = (\hat{\alpha},\ldots,\hat{\alpha})$, $\hat{\boldsymbol{\mu}}^{(0)} = h(\hat{\boldsymbol{\eta}}^{(0)})$.

(iii) Cycle: $k = 1, 2, \ldots$,

(a) Estimate all components: For $j = 1,\ldots,p$ compute

$$f_{j,\text{new}} = \mathbf{B}_j(\mathbf{B}_j^T\hat{\mathbf{W}}\mathbf{B} + \boldsymbol{\lambda}\mathbf{P})^{-1}(\mathbf{B}_j^T\hat{\mathbf{W}}\hat{\mathbf{D}}^{-1}(\mathbf{y} - \hat{\boldsymbol{\mu}})),$$

where $\mathbf{B}_j^T = (\mathbf{B}_{1j},\ldots,\mathbf{B}_{nj})$, $\mathbf{B}_{ij} = (B_{1,m}(x_{ij}),\ldots,B_{K,m}(x_{ij}))$, $\hat{\mathbf{D}} = \text{diag}(h'(\hat{\eta}_1^{(k)}),\ldots,h'(\hat{\eta}_n^{(k)}))$, $\hat{\mathbf{W}} = \hat{\mathbf{D}}\Sigma\hat{\mathbf{D}}$, and $\Sigma = \text{diag}(\text{var}(y_1),\ldots,\text{var}(y_1))$.

(b) Selection: Set $f_j = f_j^{(l)} + f_{j,\text{new}}$ yielding $\hat{\boldsymbol{\eta}}_{j,\text{new}}$ and get $s = \arg\max_j\{\text{Dev}(\hat{\boldsymbol{\eta}}^{(k)}) - \text{Dev}(\hat{\boldsymbol{\eta}}_{j,\text{new}})\}$, where $\text{Dev}(\boldsymbol{\eta})$ denotes the deviance when $\boldsymbol{\eta}$ is the fitted predictor.

(c) Update: Set $\hat{f}_{j+1}^{(k+1)} = \hat{f}_j^{(k)} + f_{s,\text{new}}$.

The R package, GAMBoost is developed based on this algorithm.

References

References

Bache, K., Lichman, M., 2013. UCI machine learning repository.
URL `http://archive.ics.uci.edu/ml`

Binder, H., Tutz, G., 2008. A comparison of methods for the fitting of generalized additive models. Stat. Comput 18, 87 – 99.

Cristianini, N., Shawe-Taylor, J., 2000. An Introduction to Support Vector Machines and Other Kernel-based Learning Methods. Cambridge University Press, The Edinburgh Building, Cambridge CB2 8RU, UK.

Eguchi, S., Copas, J., 2002. A class of logistic-type discriminant functions. Biometrika 89 (1), 1–22.

Friedman, J., Hastie, T., Tibshirani, R., 2000. Additive logistic regression: a statistical view of boosting. Ann. Statist., 337 – 407.

Hastie, T., Tibshirani, R., 1986. Generalized additive models. Statistical Science 1 (3), 297 – 318.

Hastie, T., Tibshirani, R., 1990. Generalized additive models. Vol. 43. CRC Press.

Hastie, T., Tibshirani, R., Friedman, J., 2001. The Elements of Statistical Learning. Springer, Berlin.

Little, M. A., McSharry, P. E., Roberts, S. J., Costello, D. A., Moroz, I. M., et al., 2007. Exploiting nonlinear recurrence and fractal scaling properties for voice disorder detection. BioMedical Engineering OnLine 6 (1), 23.

Marra, G., Wood, S., 2012. Coverage properties of confidence intervals for generalized additive model components. Scandinavian Journal of Statistics 39 (1), 53 –74.

Marx, B. D., Eilers, P. H., 1998. Direct generalized additive modeling with penalized likelihood. Computational Statistics & Data Analysis 28 (2), 193–209.

Minka, T., March 26 2007. A comparison of numerical optimizers for logistic regression.
URL `http://research.microsoft.com/en-us/um/people/minka/papers/logreg/`

Pepe, M. S., 2004. The statistical evaluation of medical tests for classification and prediction. Oxford University Press, USA.

Stone, C. J., 1985. Additive regression and other nonparametric models. The annals of Statistics, 689–705.

Tian, L., Cai, T., Goetghebeur, E., Wei, L. J., 2007. Model evaluation based on the sampling distribution of estimated absolute prediction error. Biometrika 94 (2), 297–311.

Tutz, G., Binder, H., 2006. Generalized additive modeling with implicit variable selection by likelihood-based boosting. Biometrics 62 (4), 961–971.

Wood, S., 2004. Stable and efficient multiple smoothing parameter estimation for generalized additive models. Journal of the American Statistical Association 99 (467).

Wood, S., 2006. Generalized Additive Models: An Introduction with R. Chapman & Hall.

Wood, S., 2008. Fast stable direct fitting and smoothness selectionfor generalized additive models. Journal of Royal Statist. Soc. B. 70 (3), 495 – 518.